

De-aliasing Undersampled Volume Images for Visualization

Galen G. Gornowicz², David H. Laidlaw¹, Jerry W. Shan², Daniel B. Lang², Paul E. Dimotakis²

¹Computer Graphics Lab

²Graduate Aeronautical Laboratory

California Institute of Technology

Pasadena, CA 91125

Abstract

We present and illustrate a new technique, Image Correlation Supersampling (ICS), for resampling volume data that are undersampled in one dimension. The resulting data satisfies the sampling theorem, and, therefore, many visualization algorithms that assume the theorem is satisfied can be applied to the data. Without the supersampling the visualization algorithms create artifacts due to aliasing.

The assumptions made in developing the algorithm are often satisfied by data that is undersampled temporally. Through this supersampling we can completely characterize phenomena with measurements at a coarser temporal sampling rate than would otherwise be necessary. This can save acquisition time and storage space, permit the study of faster phenomena, and allow their study without introducing aliasing artifacts.

The resampling technique relies on *a priori* knowledge of the measured phenomenon, and applies, in particular, to scalar concentration measurements of fluid flow. Because of the characteristics of fluid flow, an image deformation that takes each slice image to the next can be used to calculate intermediate slice images at arbitrarily fine spacing. We determine the deformation with an automatic, multi-resolution algorithm.

1 Introduction

Many visualization tools assume that they can interpolate samples in an image or volume to create a continuous function over a region. Examples include marching cubes [Lorenson and Cline, 1987], which tri-linearly interpolates sample values to find isosurfaces, and standard image supersampling [Foley et al., 1990], which uses various interpolation techniques to calculate pixel values between known ones.

If the assumption that samples can be interpolated is not valid, the visualization algorithms will produce results with artifacts. The assumption is equivalent to assuming that sampled data satisfies the sampling theorem [Oppenheim et al., 1983]. The sampling theorem states that regularly spaced point samples of a function can be used to exactly reconstruct the function if and only if the function has no content above half the sampling frequency. While some types of data satisfy the sampling theorem, others do not. When a function is not properly sampled the samples cannot, in general, be interpolated to reconstruct it. Figure 1 shows an example where artifacts are introduced by interpolative supersampling of an undersampled image. High frequency information encoded in the sample values aliases into lower-frequency artifacts in the supersampled data.

1.1 Creating Properly Sampled Data: Image Correlation Supersampling

We have developed a technique for converting certain sampled data that do not satisfy the sampling theorem into sampled data that do. Our primary application involves temporally supersampling time-varying images of flowing fluid where the temporal sampling rate is too low, but the technique may also be applied to volume data where: 1) one dimension is undersampled, 2) a parameterized deformation characterizes the change from one slice image to the next, and 3) the parameters of the deformation between two slices can be found.

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 2006		2. REPORT TYPE		3. DATES COVERED 00-00-2006 to 00-00-2006	
4. TITLE AND SUBTITLE De-aliasing Undersampled Volume Images for Visualization				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Office of Scientific Research, 875 North Randolph Street Suite 325, Arlington, VA, 22203-1768				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 11	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			




Figure 1: An example of the impact of interpolative methods for supersampling undersampled data. (a) A properly sampled image of a shape with a Gaussian cross section. (b) shows supersampling of an undersampled version of the same object. The undersampling and supersampling are in the vertical direction. Note the bumpy shapes introduced.

The algorithm we apply, which is described in Section 2 and in more detail in [Gornowicz, 1997], uses a multi-resolution deformation to characterize the change from one slice image to the next and automatically optimizes the deformation parameters at each level of resolution to find the appropriate deformation. Once the deformation has been found, it is used to calculate intermediate images at a sufficient sampling rate to satisfy the sampling theorem.

1.2 Related Work

This work has been motivated in part by metamorphosis, or morphing, techniques for both images [Beier and Neely, 1992] and volumes [Lerios et al., 1995]. Both of these techniques use carefully crafted user input to define a deformation that takes one image or volume into another. For intermediate frames they both dissolve from one partially deformed image or volume into the second. Our technique differs primarily in that the deformation we calculate is physically motivated by the equations of fluid flow, and is found automatically rather than being interactively specified.

A similar deformation-based technique [Ruprecht and Müller, 1994] defines correspondence points on different magnetic resonance images and calculates a deformation by attempting to deform the images to satisfy the correspondences. Our deformation has a more physical basis, and is parameterized and discovered differently.

Video compression techniques solve a similar problem. By using a predictive model of the deformation of one video frame to the next they can encode only the differences between a deformed version of an earlier frame and the current frame [Huang et al., 1996]. If the deformation is accurate, then the differences will be small and the resulting compression high. Because a failure in the deformation leads only to less effective compression, these techniques do not fail at representing the information if they are not robust, their compression rates degrade. Although our technique is limited to certain types of image data, it robustly supersamples them to satisfy the sampling theorem, while the video compression techniques would not.

This work is based on Image Correlation Velocimetry (ICV) [Tokumaru and Dimotakis, 1995], an automatic technique for extracting velocity information from a series of images of concentration data. That original work was subsequently extended with a multi-resolution basis for the deformation [Gornowicz, 1997]. The velocity information serves as the deformation in the two-step metamorphosis process.

Other related techniques for the computational analysis of motion from a sequence of images include particle image velocimetry [Adrian, 1991] [Willert and Gharib, 1991], variational approaches [Zhou et al., 1995] [Szeliski and Shum, 1996], and differential techniques [Horn and Schunck, 1981] [Willick and Yang, 1991] [Dahm et al., 1992]). A review of so-called optical flow techniques is given by [Barron et al., 1994].

1.3 Roadmap

Section 2 describes the algorithm and is followed by a description of the results produced by applying the algorithm to both simulated and acquired data. A discussion of the algorithm and results is in Section 5, with conclusions immediately after.

2 Image Correlation Supersampling

The Image Correlation Supersampling (ICS) algorithm consists of two parts, as do many other morphing techniques. In the first part a deformation or displacement field is defined that transforms one image into the next. In the second part, a weighted sum of the partially displaced images is used to generate intermediate images at a sampling rate sufficient to satisfy the sampling theorem. This rate can be determined from knowledge of the phenomena being measured or by estimating the rate conservatively and examining the resulting samples to calculate a more accurate and lower rate. For fluid experiments, a maximum velocity is frequently available and can be used to bound the maximum sampling rate.

In this section we first present the physically-based motivation for ICS, and then discuss the two parts of the algorithm.

2.1 Physical Motivation for Deformation

For 2D flows where a) diffusive effects are negligible (e.g. high Schmidt number or short time periods), and b) the concentration is proportional to image intensity, the scalar transport equation applies to the intensity image $I(x, t)$, and

$$\frac{\partial}{\partial t}I + u \cdot \frac{\partial}{\partial x}I \simeq 0 \quad (1)$$

where u is the velocity [Batchelor, 1967]. With τ a small step in time, and with the definition of the displacement field $\xi \equiv \tau u$, then

$$\tau \frac{\partial}{\partial t}I + \xi \cdot \frac{\partial}{\partial x}I \simeq 0 \quad (2)$$

For two concentration images, $I(x, t)$ and $I(x, t + \tau)$, $I(x, t + \tau)$ can be written as a Taylor expansion around x and t .

$$I(x + \xi, t + \tau) = I(x, t) + \tau \frac{\partial}{\partial t}I(x, t) + \xi \cdot \frac{\partial}{\partial x}I(x, t) + \text{H.O.T.} \quad (3)$$

For small ξ and τ , the higher order terms can be neglected. The sum of the first order terms in ξ and τ are also negligible from Equation 2.

2.2 Image Deformation

ICS seeks the displacement field, $\xi(x)$, that transforms one images into the next. We define $I_1(x) = I(x, t + \tau)$ and $I_0(x) = I(x, t)$.

$$I_1(x + \xi(x)) \simeq I_0(x) \quad (4)$$

$\xi(x)$ is a 2-vector field over the spatial coordinates, x , where our image is defined. $\xi(x)$ represents a displacement, or deformation, of the spatial coordinates, and hence of the image. Given the two scalar field images $I_0(x)$ at time t and $I_1(x)$ at time $t + \tau$, ICS seeks a displacement field $\xi(x)$ that minimizes the cost function $J(\xi)$ given by the square of the differences in intensity integrated over the correlation domain, Ω .

$$\min_{\xi} J(\xi) = \min_{\xi} \int_{\Omega} [I_1(x + \xi) - I_0(x)]^2 d\Omega \quad (5)$$

The displacement field, $\xi(x)$, is implemented using a multi-resolution hierarchy of tri-cubic B-splines, and the minimization of $J(\xi)$ is implemented with a multi-dimensional gradient approach. More details are available in [Gornowicz, 1997].

2.3 Supersampling

Intermediate slice images, subscripted by $0 \leq \alpha \leq 1$, are computed from the surrounding images and the displacement field, $\xi(x)$, using

$$I_{\alpha}(x) = (1 - \alpha)I_0(x - \alpha\xi) + \alpha I_1(x + (1 - \alpha)\xi) \quad (6)$$

where α varies from 0 to 1 as the image transitions from I_0 to I_1 . The equation linearly blends image I_0 to I_1 , with only image I_0 contributing when $\alpha = 0.0$ and only image I_1 contributing when $\alpha = 1.0$. Image I_0 is linearly displaced; it is undeformed when $\alpha = 0.0$ and fully displaced when $\alpha = 1.0$. I_1 is linearly inversely displaced; it is fully displaced when $\alpha = 0.0$ and undeformed when $\alpha = 1.0$.

3 Imaging

In this section we describe the two volume images created and the one acquired to test and illustrate our supersampling technique. In the two simulated volumes the sampling theorem is satisfied in two of the three directions. In the third direction the sampling theorem is not satisfied—the sampling kernel, or point spread function, is too narrow; equivalently, the samples are spaced too far apart. However, the behavior of the simulated phenomena in the regions between slices is well characterized as a deformation of the nearby slices, and thus satisfies our assumptions. In the acquired data the two spatial sampling rates are again sufficient to satisfy the sampling theorem, but the temporal sampling rate is not.

3.1 Translated Disk

For the first simulated dataset a 2D normal distribution translates from the lower left toward the upper right of the image (see Figure 2). The dataset consists only of the two images, and the deformation is a simple translation.

3.2 Translated, Rotated, Strained Images

For the second simulated dataset an image of two regions of band-limited noise are translated, rotated, and strained in two dimensions. Figure 3 shows three frames from the dataset. These deformations are exact solutions to the Navier-Stokes equations for 2D flow.




Figure 2: Frames from simulated dataset of translating disk.




Figure 3: Frames from a simulated 2D flow experiment.




Figure 4: Frames from measurements of turbulent fluid flow. Time increase for images toward the right. The overall flow within each image is toward the left.

3.3 Turbulent Jet Data

The experimentally measured data, shown in Figure 4, consists of images of a liquid-phase, axisymmetric, turbulent jet, with $Re = 2250$, discharging into a large quiescent reservoir of square cross-section. The flow was established by pressurizing the jet plenum in a blow-down, nearly-constant-pressure configuration from two gas cylinders. The jet nozzle had an exit diameter of 2.54 mm, and the discharge tank had internal dimensions of $0.85 \times 0.85 \times 1.6\text{m}$ ($335 \times 335 \times 625$ diameters). Rectangular glass windows on all four sides of the tank, in combination with a 10 inch diameter glass window at the bottom of the tank, provide optical access to the flow facility.

A laser sheet containing the axis of the jet was formed by passing the beam from a frequency-doubled Nd:YAG laser pulsed at 10 Hz through a -6.4mm focal length cylindrical lens. A long focal length ($f = +1.0\text{m}$) cylindrical lens oriented 90 degrees with respect to the first lens ensured that the laser sheet




Figure 5: 2D slices of simulated 3D data, with the undersampled axis vertical and one of the sufficiently sampled axes horizontal. (a) shows Image Correlation Supersampling (ICS), and (b) interpolative super-sampling. In the 3D data a bright disk translates from left to right, as is reconstructed correctly by ICS in (a). The disk is incorrectly split into two features in (b).




Figure 6: (a) Image Correlation Supersampling (ICS), and (b) interpolative super-sampling. The images show an isosurface calculated from the 3D data. In the 3D data a bright disk translates from left to right, as is reconstructed correctly by ICS in (a). The disk is incorrectly split into two features in (b).

reached its waist, and remained thin, in the region of interest. The 532 nm line from the laser was used to excite a fluorescent dye that was pre-mixed in weak concentration with the jet plenum fluid. The laser-induced fluorescence field was recorded at a rate of 10 frames/second on a 12-bit camera that utilizes a 1024×1024 pixel Cassini CCD. A sequence of 42 images of a 125×125 diameter region centered on the jet axis at a downstream distance of 225 diameters from the jet nozzle was recorded. For this experiment, the in-plane spatial resolution and the out-of-plane spatial (laser-sheet thickness) resolution was sufficient to resolve the diffusion scale.




Figure 7: 2D slices of 3D data measuring turbulent flow, with the undersampled temporal axis vertical and one sufficiently sampled spatial axis horizontal. (a) Image Correlation Supersampling (ICS), and (b) interpolative supersampling. (b) shows significant aliasing introduced by the supersampling. The aliasing is not apparent in (a). (c) is a gold standard for the supersampled data.




Figure 8: An isosurface calculated from simulated 2D fluid flow. (a) Image Correlation Supersampling (ICS), and (b) interpolative supersampling. (c) shows an isosurface calculated from data correctly sampled at high resolution.

4 Results

The image correlation supersampling technique has been applied to the datasets described above. We show the results of some common visualization techniques applied to the data both with and without ICS supersampling.

The first example of Image Correlation Supersampling is shown in Figures 5 and 6. The original 2D image data for this example represent a translating disk shape and are shown in Figure 2. Figure 5 shows 2D slices taken perpendicular to the original images, with one spatial dimension on the horizontal axis and time on the vertical axis. Supersampling was accomplished via ICS in 5(a) and via interpolative supersampling in 5(b). Note that ICS supersampling correctly maintains the integrity of the feature translating through the 3D volume, whereas the feature is split into two parts in (b). Figure 6 shows an isosurface created from the same data. In 6(a), the correct translation of the disk leads to a rod, as expected, while in 6(b) the rod is incorrectly divided into two separate features.

The second example of ICS is shown in Figures 7 and 8. The original 2D image data for this example represent two deforming regions of band-limited noise (see Figure 3). Figure 7 shows 2D slices taken perpendicular to the original images, with one spatial dimension on the horizontal axis and time on the vertical axis. Supersampling was accomplished via ICS in 7(a) and via interpolative supersampling in 7(b). Note that ICS supersampling correctly maintains the integrity of features translating through the 3D volume, whereas the features are distorted in 7(b). Figure 7(c) shows the data directly sampled at high temporal




Figure 9: An isosurface calculated from 3D imaging data of turbulent fluid flow. (a) Image Correlation Supersampling (ICS), and (b) interpolative supersampling.



Figure 10: (a) Image Correlation Supersampling (ICS), and (b) interpolative supersampling. The images show 2D slices of 3D data measuring turbulent flow. (b) shows significant aliasing introduced by the supersampling. The aliasing is not apparent in (a).

resolution from the analytic expression for the fluid flow. It is a gold standard for the supersampled data. Figure 8 shows an isosurface created from the same data. In 8(a), the correct deformation of the noise images leads to two slightly twisted tubes, as expected, while in 8(b) the tubes are split into pieces and are significantly distorted. 8(c), a gold standard for this test case, shows what the isosurface should look like given perfect sampling.

We have also applied the ICS algorithm to data imaged from turbulent flow. The data for this example are shown in Figure 4, and are sampled at 10 frames/second. The results are supersampled at 200 frames/second, a rate determined to be sufficient from the velocity of the jet. The 45° angle of features in Figure 10 further suggests that the rate is sufficient because the temporal sampling of the image is similar to the spatial sampling, where the sampling rate is known to be sufficient.

Figure 9(a) shows an isosurface of concentration data measured in a plane. Time is stacked in the third dimension. The jet flows towards the left, and time moves from top to bottom in the box. Note that structures moving through time trace out smooth, slanted surfaces, with steep surfaces representing slow structures, and shallow surfaces fast structures. In Figure 9(b) interpolative supersampling introduces many artifacts,

analogous to those shown in earlier figures, which are incorrect and which mislead the viewer about the behavior of the fluid. Note the ridges on the left-most surface in the image, and the three separate shapes at the far right corresponding to a single rod in Figure 9(a). Figure 10 shows a 2D section of the 3D data. Note the artifacts in 10(b) created by aliasing. They are absent in 10(a), the ICS image.

5 Discussion

We have made some observations about ICS and its application while working to apply it to the example datasets.

5.1 Interpreting Image Displacements as Velocity

The ICV algorithm [Tokumaru and Dimotakis, 1995] interprets the displacement field that it calculates as a velocity. For some cases, in particular where image intensities are constant within a region, the velocity is unconstrained, and so there are many possible solutions. This ambiguity leads to uncertainty in the velocities. For ICS, however, the ambiguity in the displacement is canceled out in the supersampling step because the displacement ambiguity lies completely within regions of constant intensity.

5.2 Out-of-plane Motion

The physical basis for using a displacement field to represent fluid behavior between images is based on 2D flow within the images. In cases where there is flow out of the imaging plane the changes are not correctly accounted for in intermediate images. Some artifacts, then, remain in the supersampled data.

The jet data that we measured (Figure 4) does have fluid flow components out of the imaging plane. Because of the orientation of the jet along the imaging plane, the out-of-plane velocity components are smaller than $\frac{1}{3}$ the in-plane components, and so are not as important. There are also significant portions of the dataset where the out-of-plane velocities are negligible. Even with this out-of-plane velocity, ICS reduces the artifacts due to high-frequency in-plane velocities, and does not increase the artifacts in regions of out-of-plane motion. We feel that the costs, then, are minimal and the benefits significant, although any artifacts that remain must be carefully interpreted.

To address this problem we propose two possibilities. The first is higher-speed acquisitions, although the associated cost is high. The second is collecting volume data, so that a displacement field can be calculated that takes into account the 3D flow.

5.3 Time-sampled Volume Data: the Fourth Dimension

A logical extension of this algorithm would apply to 4D data—volumes that are sampled over time. The deformation or displacement, ξ , would be extended to 3D. This extension would slow the optimization process, but would be particularly useful in this context because it would help reduce the temporal sampling rate and hence reduce the long acquisition times and large data size for these datasets.

5.4 Other Application Areas

With an extension to volumes, and an appropriate modification to the type of displacement or deformation allowed, this work should apply to several medical imaging areas. MRI time-series are often undersampled temporally, and so do not sample the motion of the subject well. The body, however, is generally displacing from volume image to volume image, and the displacement should be calculatable and usable to supersample the data temporally.

Similarly, developmental studies of animals track the changes in shape of different features. If these can be modeled as displacements within volume images, with expansion, contraction, and deformation of the shapes taken into account, then a relatively small number of volume images of a developing animal might be used to reconstruct a continuous representation of the animal during development.

6 Conclusions

The technique that we have presented, and others with a similar physical motivation, can be used to reduce the high sampling rates required of acquired data to satisfy the sampling theorem by post-processing the undersampled data. Computation and visualization tools that assume the sampling theorem is satisfied can then be applied to the measurements. The savings in collection time, higher-speed hardware, and storage can be significant. The acquired dataset illustrating the algorithm would have had to be sampled 20 times as fast to achieve similar results without ICS. To collect at that rate would have required significant hardware upgrades for both the excitation laser and for the camera and data acquisition hardware—upgrades that approach and may even exceed the capabilities of available technology. With ICS, existing lower-cost hardware was sufficient, and the behavior of the jet can still be studied with standard visualization tools.

7 Acknowledgements

This work was supported in part by grants from Apple, DEC, Hewlett Packard, and IBM. Additional support was provided by NSF (ASC-89-20219) as part of the NSF/ARPA STC for Computer Graphics and Scientific Visualization, by the National Institute on Drug Abuse and the National Institute of Mental Health as part of the Human Brain Project, and by the AFOSR (F49620-94-1-0353, F49620-93-1-0338, and F49620-95-1-0199) as part of a larger effort to understand turbulent mixing and chemical reactions in free-shear flows. Thanks also to Pavel Svitek for support with mechanical design and other assistance in the experiments, and to David Breen for reviewing early drafts.

References

- [Adrian, 1991] Adrian, R. J. (1991). *Particle-Imaging Techniques for Experimental Fluid Mechanics*.
- [Barron et al., 1994] Barron, J., Fleet, D., and Beauchemin, S. (1994). Performance of optical flow techniques. *Int. J. Comp. Vision*, 12:43–77.
- [Batchelor, 1967] Batchelor, G. K. (1967). *An Introduction to Fluid Dynamics*. Cambridge University Press, Cambridge.
- [Beier and Neely, 1992] Beier, T. and Neely, S. (1992). Feature-based image metamorphosis. In Catmull, E. E., editor, *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 35–42.
- [Dahm et al., 1992] Dahm, W. J. A., Su, C. C., and Southerland, K. B. (1992). A scalar imaging velocimetry technique for fully resolved four-dimensional vector velocity measurements in turbulent flows. *Physics of Fluids A*, 4(10):2191–2206.
- [Foley et al., 1990] Foley, J. D., van Dam, A., Feiner, S. K., and Hughes, J. F. (1990). *Computer Graphics, Principles and Practice, Second Edition*. Addison-Wesley, Reading, Massachusetts. Overview of research to date.

- [Gornowicz, 1997] Gornowicz, G. G. (1997). *Continuous-Field Image-Correlation Velocimetry And Its Application To Unsteady Flow Over An Airfoil*. Eng thesis, California Institute of Technology.
- [Horn and Schunck, 1981] Horn, B. K. P. and Schunck, B. G. (1981). Determining optical flow. *Artificial Intelligence*, 17:185–203.
- [Huang et al., 1996] Huang, T. S., Stroming, J. W., Kang, Y., and Lopez, R. (1996). Advances in very low bit rate video coding in north america. *IEICE Transactions on Communications*, E79-B(10):1425–1433.
- [Lerios et al., 1995] Lerios, A., Garfinkle, C. D., and Levoy, M. (1995). Feature-Based volume metamorphosis. In Cook, R., editor, *SIGGRAPH 95 Conference Proceedings*, Annual Conference Series, pages 449–456. ACM SIGGRAPH, Addison Wesley. held in Los Angeles, California, 06-11 August 1995.
- [Lorenson and Cline, 1987] Lorenson, W. E. and Cline, H. E. (1987). Marching cubes: A high resolution 3D surface construction algorithm. In Stone, M. C., editor, *Computer Graphics (SIGGRAPH '87 Proceedings)*, volume 21, pages 163–169.
- [Oppenheim et al., 1983] Oppenheim, A. V., Willsky, A. S., and Young, I. T. (1983). *Signals and Systems*. Prentice-Hall, Inc., New Jersey.
- [Ruprecht and Müller, 1994] Ruprecht, D. and Müller, H. (1994). Deformed cross-Dissolves for image interpolation in scientific visualization. *The Journal of Visualization and Computer Animation*, 5(3).
- [Szeliski and Shum, 1996] Szeliski, R. and Shum, H. (1996). Motion estimation with quadtree splines. *IEEE Trans. Pattern Matching and Machine Intelligence*, 18(12):1199–1210.
- [Tokumaru and Dimotakis, 1995] Tokumaru, P. T. and Dimotakis, P. E. (1995). Image correlation velocimetry. *Experiments in Fluids*, 19(1):1–15.
- [Willert and Gharib, 1991] Willert, C. E. and Gharib, M. (1991). Digital particle image velocimetry. *Experiments in Fluids*, 10:181–193.
- [Willick and Yang, 1991] Willick, D. and Yang, Y. H. (1991). Experimental evaluation of motion constraint equations. *CVGIP: Image Understanding*, 54(2):206–214.
- [Zhou et al., 1995] Zhou, Z., Synolakis, C. E., Leahy, R. M., and Song, S. M. (1995). Calculation of 3d internal displacement fields from 3d x-ray computer tomographic images. *Proc. Roy. Soc. London A*, 449:537–554.